CLAIMS

1. A computer-implemented anomaly detection method comprising employing at least one hardware processor of a computer system, in response to selecting a training token sequence from a training corpus of token sequences and in response to selecting a transformation from a pre-determined plurality of sequence transformations, to distort the training token sequence by applying the selected transformation to the training token sequence to produce a modified token sequence, the method characterized by further employing the at least one hardware processor to:

input the modified training sequence to a sequence analyzer comprising:

a first prediction head configured to determine a first prediction indicator according to the modified token sequence, the first prediction indicator indicative of a likelihood that the selected transformation was applied to produce the modified token sequence, and

a second prediction head configured to determine a second prediction indicator according to the modified token sequence, the second prediction indicator indicative of a likelihood that a selected token of the training token sequence was changed by the application of the selected transformation;

in response to the sequence analyzer determining the first and second prediction indicators, adjust at least one parameter of the first prediction head according to the second prediction indicator and adjust at least one parameter of the second prediction head according to the first prediction indicator; and

in response to adjusting parameters of the first and second prediction heads, employ the sequence analyzer to determine whether a target token sequence is anomalous.

2. The method of claim 1, wherein applying the selected transformation comprises replacing the selected token of the training token sequence with a substitute token.

3. The method of claim 2, further characterized by employing the at least one hardware processor to:

execute a token generator configured to generate the substitute token according to the training token sequence; and

in response to the sequence analyzer determining the first and second prediction indicators, adjust at least one parameter of the token generator according to the first or second prediction indicator.

4. The method of claim 1, wherein applying the selected transformation comprises an item selected from a group consisting of deleting the selected token from the training token sequence, inserting an additional token into the training token sequence, and permuting a selected subset of tokens of the training token sequence.

5. The method of claim 1, wherein:

the training and target token sequences comprise text formulated in a natural language; and

the method is further characterized by having the at least one hardware processor, in response to determining that the target token sequence is anomalous, further determine that an author of the target token sequence differs from an author of the training token sequence.

6. The method of claim 1, wherein:

the training and target token sequences comprise text formulated in a natural language; and

the method is further characterized by having the at least one hardware processor, in response to determining that the target token sequence is anomalous, further determine that a subject matter of the target token sequence differs from a subject matter of the training token sequence.

7. The method of claim 1, wherein:

the training and target token sequences comprise text formulated in a natural language; and

the method is further characterized by having the at least one hardware processor, in response to determining that the target token sequence is anomalous, further determine that the target token sequence is machine-generated.

8. The method of claim 1, wherein:

the training and target token sequences consist of sequences of computing events; and;

the method is further characterized by having the at least one hardware processor, in response to determining that the target token sequence is anomalous, further determine that the target token sequence is indicative of a computer security threat.

9. A computer system comprising at least one hardware processor configured to execute an input modifier and a sequence analyzer, wherein the input modifier is configured, in response to selecting a training token sequence from a training corpus of token sequences and in response to selecting a transformation from a pre-determined plurality of sequence transformations, to distort the training token sequence by applying the selected transformation to the training token sequence to produce a modified token sequence, the computer system characterized by further being configured to:

input the modified token sequence to the sequence analyzer, wherein the sequence analyzer comprises:

a first prediction head configured to determine a first prediction indicator according to the modified token sequence, the first prediction indicator indicative of a likelihood that the selected transformation was applied to produce the modified token sequence;

a second prediction head configured to determine a second prediction indicator according to the modified token sequence, the second prediction indicator indicative of a likelihood that a selected token of the training token sequence was changed by the application of the selected transformation;

in response to the sequence analyzer determining the first and second prediction indicators, adjust at least one parameter of the first prediction head according to the second

prediction indicator and adjust at least one parameter of the second prediction head according to the first prediction indicator; and

in response to adjusting parameters of the first and second prediction heads, execute the sequence analyzer to determine whether a target token sequence is anomalous.

10. The computer system of claim 9, wherein applying the selected transformation comprises replacing the selected token of the training token sequence with a substitute token.

11. The computer system of claim 10, further characterized by being further configured to:

execute a token generator configured to generate the substitute token according to the training token sequence; and

in response to the sequence analyzer determining the first and second prediction indicators, adjust at least one parameter of the token generator according to the first or second prediction indicators.

12. The computer system of claim 9, wherein applying the selected transformation comprises an item selected from a group consisting of deleting the selected token from the training token sequence, inserting an additional token into the training token sequence, and permuting a selected subset of tokens of the training token sequence.

13. The computer system of claim 9, wherein:

the training and target token sequences comprise text formulated in a natural language; and

the computer system is further characterized by being further configured to, in response to determining that the target token sequence is anomalous, determine that an author of the target token sequence differs from an author of the training token sequence.

14. The computer system of claim 9, wherein:

the training and target token sequences comprise text formulated in a natural language; and

the computer system is further characterized by being further configured to, in response to determining that the target token sequence is anomalous, determine that a subject

matter of the target token sequence differs from a subject matter of the training token sequence.

15. The computer system of claim 9, wherein:

the training and target token sequences comprise text formulated in a natural language; and

the computer system is further characterized by being further configured to, in response to determining that the target token sequence is anomalous, determine that the target token sequence is machine-generated.

16. The computer system of claim 9, wherein:

the training and target token sequences consist of sequences of computing events; and;

the computer system is further characterized by being further configured to, in response to determining that the target token sequence is anomalous, determine that the target token sequence is indicative of a computer security threat.

17. A non-transitory computer-readable medium storing instructions which, when executed by at least one hardware processor of a computer system, cause the computer system, in response to selecting a training token sequence from a training corpus of token sequences and in response to selecting a transformation from a pre-determined plurality of sequence transformations, to distort the training token sequence by applying the selected transformation to the training token sequence to produce a modified token sequence, the computer-readable medium characterized by further storing instructions causing the computer system to:

input the modified token sequence to a sequence analyzer comprising:

a first prediction head configured to determine a first prediction indicator according to the modified token sequence, the first prediction indicator indicative of a likelihood that the selected transformation was applied to produce the modified token sequence, and

a second prediction head configured to determine a second prediction indicator according to the modifed token sequence, the second prediction indicator indicative of a likelihood that a selected token of the training token sequence was changed by the application of the selected transformation;

in response to the sequence analyzer determining the first and second prediction indicators, adjust at least one parameter of the first prediction head according to the second prediction indicator and adjust at least one parameter of the second prediction head according to the first prediction indicator; and

in response to adjusting parameters of the first and second prediction heads, execute the sequence analyzer to determine whether a target token sequence is anomalous.